

## Use of a Job Ticket as a Generic XML Database

### Technical Field

The technical field is the integration and control of services in a networked environment.

### Background

Services may be provided by one or more operating units in a computer-based network. Users of the network may generate specific tasks and send the tasks into the network to be assigned to one of the operating units. For example, a user at a computer terminal may generate a printing order using a printer driver installed on the terminal. The printer driver is used to control the printing request. In another example, a user at a computer terminal may generate a printing order and send the printing order into a computer network so that the printing order is completed by a printing service. The printing order may be related to a company brochure. The printing order may contain unique requirements such as paper type, font size, layout, graphics, color, and other requirements. The user may specify that a specific printing service, such as Kinkos, prepare the company brochure. Alternatively, the computer network may include programs that suggest printing services to the user.

To control the printing job, the user's computer terminal may generate a job ticket. The job ticket includes the requirements, such as the requirements listed above, and an identification of the specific job that allows the job status to be tracked through the computer network.

Use of the job ticket allows printing and similar services to be allocated to those resources (i.e., the operating units) that are best suited to completing the services. Unfortunately, current computer systems do not allow access to the wide variety of services existing in networked computer systems, such as the Internet. In addition, current systems require users to have some knowledge of the existing resources, and may require users to include applicable programming to communicate with the services. Furthermore, current systems do not allow a job request to be split among several processors. As a result, completion of the job request may take longer than necessary, and may not be completed in the most efficient, lowest-cost manner.

### Summary

To overcome these and other problems related to use of a job ticket, a method and an apparatus allow a client to manage job attributes and processes using an electronic service center. The service center includes a job ticket service that allows access and modification of a job ticket

00873192:060501  
105090

by multiple users on a network. The method and apparatus use a network-accessible job ticket to relate to a specific job or content. The job ticket may be an object, such as an XML object, comprising routines and data. The content may be stored on the network and may be accessed by multiple job tickets. Storage and management of the job ticket are transparent to the user. The job ticket is stored in a common location in the network. The job ticket remains in the same location in the network, and users access only that portion of the job ticket required to complete a designated process. Security measures may be added to limit access to those users designated as being allowed to access the job ticket and the job file. The job ticket may include a service ID that relates the job ticket back to the originating job ticket service. In this way, a user who acquires all or part of the job ticket can refer back to the originating job ticket service (and the original, or as-modified, job ticket) to verify any changes and to ensure that the job ticket being accessed is up-to-date. The job ticket also includes a job ID to refer the job ticket to a specific job.

The service center is coupled through a communications network to a front end service. The front end service allows a user to generate a service or job request. The communications network may be the Internet, or a local area network, for example.

The service center includes a service bus, to which are coupled a job store, the job ticket service, and a work flow controller. Also coupled to the service center are one or more processors that may be controlled to complete processes and tasks defined in the job tickets.

The job ticket service may generate and store the job tickets. Job content (e.g., a PDF file) is stored in the job store. With this structure, the user does not have to manage storage of the job content or to know which job store holds the job content. The job ticket service controls access to the job tickets, and, through the use of the job tickets, also controls access to job content in the job store, or elsewhere in the network. The job ticket service may create a reference to the job ticket, and may use the reference to control access to the job ticket.

The use of job tickets allows clients to define databases, and to store data through the job ticket service. The databases may be used to hold contact lists, addresses, and other personal data. The databases may also be used to store any other generic data. The databases could then be used in conjunction with a variety of e-services provided by the processors. For example, an e-mail processor that provides e-mail services may be used in conjunction with a personal contact list to send e-mail messages, transfer electronic files, or to establish a chat room. The e-mail

processor may access the contact list at predefined intervals to send e-mail messages to a select group of e-mail addressees. Furthermore, because the service center provides a single portal to processors that are coupled to the communications network, the client need not have any knowledge of the database structure, or the processing requirements of the processors.

In an embodiment, the job ticket may be used as a generic database in conjunction with an e-mail service. A client may have established a list of e-mail contacts. The contacts database may then be stored in the job store. A corresponding job ticket may be stored at the job ticket service. The job ticket includes control data needed to send and receive e-mail through the service center. In an embodiment, the generic database may be an XML or HTML database, or any other markup language database. The database may also be a structured query language (SQL) database.

The use of the job ticket also allows for parsing, searching and updating the contacts database. For example, the client may desire to search the contacts database for phone numbers by name. This search functionality is included in the job ticket, and allows the job ticket service to provide the client with a list of phone numbers for all entries in the contacts database by name.

#### **Description of the Drawings**

The detailed description will refer to the following figures in which like numeral refer to like items, and in which:

Figure 1 is a block diagram showing a prior art use of a job ticket;

Figure 2 is a tree diagram showing the processes in an example job ticket;

Figure 3 is a block diagram of a digital image work flow network;

Figure 4 is a block diagram of a service center used with the network of Figure 3;

Figures 5A-5D illustrate an exemplary job ticket;

Figure 6 is a diagram of functions controlled by a job ticket service;

Figure 7 is a diagram showing access functions controlled by the job ticket service;

Figure 8 is a block diagram illustrating additional control features of the job ticket service;

Figure 9 is a flow chart illustrating one of the processes controlled by the job ticket service;

and

Figures 10-16 are flow charts showing sub-processes in the overall process illustrated in Figure 9.

## Detailed Description

Figure 1 is a block diagram showing a prior art application of a job ticket service. Job tickets are often associated with a printing standard, the job definition format (JDF). The JDF is described in detail in JDF Specification Draft Spiral 4.0, available at [www.hp\\_opensource.com](http://www.hp_opensource.com), which is hereby incorporated by reference. In Figure 1, a user 1 generates a job request and sends the job request through a portal 4 to a processor 5. The job request may include a job ticket data file 2 and a content file 3. The user 1 may be a computer terminal in a networked computer system and the processor 5 may be a networked printer. The job request may involve printing a document. The document may be represented by the content 3, which is a digital representation of text and images to be printed. The intended format of the printed document may be described in the job ticket file 2, which is simply a digital file that specifies how the printer is to print the document. For example, the job ticket file 2 may require that the document be printed on back-to-back pages.

In a specific application, the functions of the job ticket file 2 may be carried out by a printer driver. The printer driver encodes control data related to printing the document, and sends the control data and the content 3 to the printer (i.e., the processor 5). The printer accesses the control data and the content 3 to print the document.

While the application shown in Figure 1 works well to print a document, the application has many drawbacks. In particular, if multiple processors are involved in producing the document, each such processor will require access to the job ticket file 2. This access brings problems related to security, modification control and workflow control. For example, each processor requiring access to the job ticket file 2 may have to wait on processing until a prior processor has completed use of the job ticket file 2. Thus, the prior art application may result in unwanted delays in completing the job request.

Prior art applications of job ticket services also suffer because the user may not know anything about the processors, including capabilities and availabilities of the processors, or even if the processors exist. Thus, the user may not know which portal to use to connect to a specific processor.

These and other problems are solved by a method and an apparatus that controls access to a job ticket and associated content through use of a job ticket service. The job ticket service

1 includes mechanisms that arbitrate access to the job ticket among multiple users of the job ticket,  
2 limit access to the job ticket by incorporating security features, and ensure modifications made by  
3 one processor or user are reflected in the job ticket and the content. In effect, the apparatus  
4 includes a generic database that couples input data from clients as job requests with output services  
5 such as processors that perform tasks or processes to complete the job requests. The database  
6 may have has the features of a generic XML database in that it is extensible, and in that the clients  
7 need not have any knowledge of the individual processes to be performed, or the internal  
8 programming requirements of the processors. Thus, the clients may submit job requests to a  
9 service center that will ensure that an appropriate processor or processors are assigned to complete  
10 the job request. Other database formats may also be used.

11 Before describing the apparatus and method in detail, a review of a job ticket is provided.  
12 Figure 2 is a node-tree diagram (or simply a node tree) 10 that illustrates processes defined in a  
13 job ticket for printing a brochure. The brochure may be printed on a commercial press, and may  
14 use digital content to generate plates for printing the brochure. Within the node tree 10, the nodes  
15 specify a product, process, or group of processes. Each node may modify, consume or create  
16 resources. Each node may contain further nested nodes, or sub-nodes. The arrangement of nodes  
17 and sub-nodes may be likened to a tree, and each node and sub-node may be referred to as a  
18 branch. A brochure node 11 defines the features and parameters of the brochure. A cover node  
19 12 defines the parameters for producing the brochure cover. Inside pages node 13 includes the  
20 parameters to produce the inside pages. The inside pages node 13 is shown with several sub-  
21 nodes, including a sub-node 14 for digital plate making. The digital plate making sub-node 14 itself  
22 includes two additional sub-nodes, a ripping sub-node 16 and a plate making sub-node 18.

23 Each of the nodes and sub-nodes shown in Figure 2 has associated with it input resources  
24 and at least one output resource. A resource may be described by parameters or logical entities.  
25 The resource may be a physical entity such as a component, a handling resource, or a consumable.  
26 A component resource may be the output of a node or sub-node, such as a printed sheets. A  
27 handling resource is used during a process, but is not consumed by the process. A consumable  
28 resource may be partly or wholly consumed by the process. Examples of consumable resources  
29 include inks, plates, and glue. Other resources may be a digital file or representation of a physical  
30 object. For example, the ripping sub-node 16 may include as input resources a run list, media, RIP

1 parameters, and layout. The run list resource describes the pages, including the files in which the  
2 pages occur, and which pages are to be used. The media resource describes the media that will  
3 be used to make plates, and is needed to describe the dimensions of the media. The RIP  
4 parameters resource describes all device-specific parameters of the ripping process. The layout  
5 resource describes placement of source pages onto the plates, and eventually onto press sheets.  
6 As an output resource, the ripping sub-node 16 may provide ripped flats. Other resources include  
7 parameter resources, which define the details of processes, as well as other non-physical computer  
8 files used by a process.

9 The node tree 10 shown in Figure 2 is intended to apply to printing a document. However,  
10 node-tree diagrams may be used to represent job tickets for other services besides printing. For  
11 example, a job ticket may be used for data processing, image processing, creating and maintaining  
12 a database, electronic publishing, e-mail, and various e-commerce services. Moreover, the job  
13 ticket may be used to allow different e-commerce services to interact with each other.

14 Figure 3 is a block diagram of a digital imaging work flow (DIW) network 20 that  
15 incorporates a service center and a job ticket service to control tasks submitted by clients. The  
16 service center may operate as a single portal through which the clients connect to one or more e-  
17 services including e-mail, e-commerce and online shopping, e-printing, and data services, including  
18 database searching and database construction, population and maintenance. Using a single portal,  
19 such as the service center, allows the clients to select from a wide variety of e-services, such as  
20 those noted above, without requiring the clients to have any prior knowledge of the e-services.

21 The service center may include components that receive information in the form of job  
22 requests, and using the information, create a job ticket that specifies tasks and resources. The job  
23 ticket may be stored in a job ticket service, and notices may be posted to indicate when a job ticket  
24 is available. Processors coupled to the service center may bid on completion of the job ticket, and  
25 the service center may include a bidding service that evaluates bids. The service center may select  
26 one or more processors to assign to the job ticket based on client-supplied criteria, or based on  
27 a set of standard criteria, including industry standard criteria. The service center may provide  
28 mechanisms to control access to the job ticket, or to portions (branches) of the job ticket. The  
29 mechanisms include branch locking, and authorization and authentication servers that use public key  
30 encryption, or similar processes.

1 The service center may include hardware components such as servers, computers, central  
2 processing units, communications interfaces, and memory devices to provide the processing  
3 capability and data storage required to carry out the above-described functions.

4 The DIW network 20 includes a front end service 30 that allows a client 31 to generate  
5 and submit a service or job request. In an embodiment, the front end service 30 may be an Internet  
6 web browser. Alternatively, the front end service 30 may be a web application or a port monitor.  
7 The job request may contain detailed information about how the job is to be executed, and may  
8 be formatted according to the job definition format standard. Alternatively, the job request may  
9 include only basic information, which will be used by another component to finalize the job  
10 definition, or work flow. Finally, the job request may include the content, or job, that is to be  
11 processed. The content could be one or more digital files, text files, and other files. The front end  
12 service 30 is coupled to a communications network 35, which may be the Internet or a local area  
13 network, for example. Coupled to the communications network 35 is a service center 40 that links  
14 one or more processors 80, to the communications network 35. Each of the processors 80, may  
15 include a cache 81, that may be used to store information related to a job request, including  
16 information related to job tickets. In an embodiment, the service center 40 may be an Internet web  
17 site that includes data storage and control functions. In another embodiment, the service center 40  
18 is a node in a local area network.

19 The service center 40 allows a broad spectrum of communications between entities coupled  
20 to the service center 40. In particular, the service center 40 allows different e-services to interact  
21 programmatically with one another using specific protocols and generic protocols (e.g., TCP/IP).  
22 This programmatic interaction allows different services and processes that are coupled to the  
23 network to exchange data and files, and to modify the data and files. The programmatic interaction  
24 may be completed by use of a remote procedure call (RPC) between entities coupled to the service  
25 center 40. Other methods for providing the programmatic interaction include CORBA, UDDI, and  
26 e-speak.

27 Figure 4 is a diagram of the service center 40. The service center 40 includes a service bus  
28 41 in communication with the communications network 35 and the processors 80, of Figure 3.  
29 Coupled to the service bus 41 is a job store 50, a job ticket service 60, a work flow controller 70,  
30 an optional bidding service 90, an authorization server 92 and an authentication server 94. The job

00073422-060501

store 50 may store one or more job content files 51. The job ticket service 60 may control one or more job tickets 61. The work flow controller 70 may use one or more agents 71, to control processes on the service bus 41.

The job store 50, job ticket service 60 and work flow controller 70 function to accept information from the clients 31, and to use the information to control the actions of the processors 80. The processors 80, performs specific tasks or processes as determined by the service center 40.

The job store 50 may be a node on the service bus 41, and may include programming to allow the job store 50 to carry out its functions. The job store 50 may be used to store the content 51, which may be in the form of one or more large files. In the context of printing a document using a service or process coupled to the service bus 41, the job store 50 may store the document content in one or more PDF files, for example. The content 51 may include graphics and text. The content 51 for a specific document may include several files. For example, a brochure may have a separate file for the cover and another file for the inside pages. Text for the inside pages may be in one file and images in yet another file. The content 51 may also include links to other resources or entities on the service bus 41. The job store 50 provides for mass storage of the content 51, so that a user (client 31 or processor 80) does not have to provide the mass storage required for the job content 51. By using the mass storage capabilities of the job store 50, the content 51 may be made to persist in the network 20, and may be made accessible to users at any time. That is, the job store 50 may store the content 51 for an extended time. The job store 50 also manages and controls the content so that the user (client 31 or processor 80) does not have to manage the content 51. Management functions include maintaining configuration or version control of the content 51, controlling access to the content 51, and maintaining the content 51 in storage.

The job ticket service 60 holds job tickets 61. The job ticket service 60 controls access to and may manage configuration of the job tickets 61. For example, the job ticket service 60 may allow users (clients 31, and processors 80,) to access a portion or branch of a job ticket 61 rather than passing the job ticket 61 among multiple users. Access to the job ticket portion may be effectuated by use of an application programming interface, a scriptable interface, or a similar feature. As noted above, the job ticket 61 does not include the content 51 (e.g., the graphical and text files of a document), but the job ticket 61 relates to content 51 (e.g., a PDF file) stored in the



1 job store 50. The user does not have to manage storage of the job content or to know which job  
2 store 50 holds the job content. The job ticket service 60 instead passes a reference in the job  
3 ticket 61. This allows multiple clients 31 and processors 80<sub>i</sub> to access the content 51.  
4 Furthermore, the content 51 may relate to more than one job ticket 61. The job ticket service 60,  
5 and its interrelationships with other entities coupled to the service bus 41, will be described later  
6 in detail.

7 Some job tickets 61 may be accessed by multiple processors 80, in either serial,  
8 overlapping, or simultaneous fashion. The multiple access processing could result in problems with  
9 use of the job ticket 61. For example, a first processor may acquire the job ticket 61 (or a portion  
10 or branch thereof), and perform a process specified in the work flow, which may modify the  
11 branch. Such modification may be to indicate a branch as complete, use up input resources, or  
12 create new output resources, for example. A second processor could attempt to acquire the  
13 branch, but might not “know” that the first processor had modified the branch. Alternatively, if two  
14 processors compete for the same branch, a deadlock situation might occur.

15 One solution to the above problems may be to lock the job ticket 61 whenever a processor  
16 80 acquires the job ticket 61. Unfortunately, locking the job ticket 61 may prevent concurrent or  
17 parallel processing and may slow down completion of the job request 32.

18 The job ticket service 60 shown in Figure 4 overcomes these and other problems by having  
19 the capability to lock the job ticket 61 at the branch level. The branch locking may be  
20 accomplished by one of several methods. The work flow controller 70 may assign one or more  
21 specific processors 80<sub>i</sub> to perform the tasks identified with the branch to be locked. Where only  
22 one processor 80 is authorized access to the branch, branch locking may not be required. Where  
23 more than one processor 80 is authorized access to the same branch, the job ticket service 60 may  
24 lock the branch when one of the authorized processors 80 actually acquire the branch.

25 If the work flow controller 70 has not assigned processors 80<sub>i</sub> to branches (i.e., any  
26 processor 80 may access a branch at any time), the job ticket service 60 may lock the branch when  
27 a processor 80 acquires the branch.

28 The job ticket service 60 may lock the branches by setting a lock/unlock flag for each  
29 branch. Processors 80<sub>i</sub> accessing the job ticket 61 may then review the lock/unlock flag status to  
30 determine if the branch may be accessed. In some circumstances, the job ticket service 60 may

1050501  
060501  
2050501  
3050501  
4050501  
5050501  
6050501  
7050501  
8050501  
9050501  
1050501  
1150501  
1250501  
1350501  
1450501  
1550501  
1650501  
1750501  
1850501  
1950501  
2050501  
2150501  
2250501  
2350501  
2450501  
2550501  
2650501  
2750501  
2850501  
2950501  
3050501

allow access only to those branches that are unlocked. A processor 80 that has completed a task defined by the branch may need to have the branch unlocked in order to modify the branch.

The work flow controller 70 may be used to create the job tickets 61 that are stored in the job ticket service 60. The work flow controller 70 may review the job requests 32 submitted by the clients 31, and may then use a job ticket template to prepare the job ticket 61. The work flow controller 70 may then send the job ticket 61 to the job ticket service 60 for storage and processing.

The work flow controller 70 also controls completion of tasks among the processors 80. In an embodiment, the work flow controller 70 determines which of the processors 80 have the necessary and available resources to begin the processes listed in a specific job ticket 61. The work flow controller 70 then designates the appropriate processors 80<sub>i</sub> to complete the tasks referenced by the job ticket 61. For example, if a job ticket 61<sub>1</sub> requires color printing, the work flow controller 70 may determine that only processor 80<sub>3</sub> is a color printer with the capacity to begin the job specified in the job ticket 61<sub>1</sub>. This embodiment in which the work flow controller 70 determines which processors 80<sub>i</sub> to assign to a specific job ticket 61 may be especially appropriate when the network 35 is a local area network and all processors 80<sub>i</sub> are directly coupled to the local area network 35.

Alternatively, the work flow controller 70 may receive bid information from Internet-connected processors 80<sub>i</sub> and may use the bid information to select the processors 80<sub>i</sub> to complete the job request 32.

The work flow controller 70 may also be used to designate the various nodes, input and output resources, and other features of the node tree used to complete the job request. That is, the work flow controller 70 may be used to create a construct, or work flow, such as the node tree 10 shown in Figure 2. To accomplish these tasks, the work flow controller 70 may include one or more agents 71<sub>i</sub> that write a job definition file, based on control data contained in the job request 32. Alternatively, a separate management information system (not shown) may be used to create the nodes, and to control flow of tasks to the processors 80<sub>i</sub> and other entities. In yet another embodiment, the job definitions may be written by the client 31 that originated the job request 32.

Referring again to the node tree 10 of Figure 2, many output resources of the individual nodes serve as input resources for other nodes. These other nodes may not be able to begin

09673102-060504  
10509072670

executing until all input resources are complete and available, which means that the nodes may need to execute in a well-defined sequence. For example, a process for making plates will produce press plates as an output resource that is required by a printing process. In the hierarchical organization of the node tree 10, nodes that occur higher in the node tree 10 represent higher-level, more abstract operations, while lower order nodes represent more detailed, specific processes. Moreover, nodes near the top of the node tree 10 may represent only intent regarding the components or assemblies that comprise the product, and lower level nodes provided the detailed instructions to a processor 80 to perform a specific process.

Because two node trees may not be similar, the work flow controller 70 may determine processes to be completed, the order in which the processes are completed, and the processors 80<sub>i</sub> that are to complete the processes. The work flow controller 70 may use the agents 71 to determine an actual work flow, considering factors such as control abilities of the processors 80<sub>i</sub> that complete the processes, transport distances between processors 80<sub>i</sub>, load capabilities of the processors, and time constraints in the job request, for example. The agents 71 may define the overall process using serial processing, which involves subsequent production and consumption of resources by the processors 80<sub>i</sub>, overlapping processing, which involves simultaneous consumption and production of resources by more than one processor 80<sub>i</sub>, parallel processing, which involves sharing resources among processors 80<sub>i</sub>, and iterative processing, which involves a back and forth processing scheme to develop resources.

Returning to Figure 4, the processors 80<sub>i</sub> may be used to provide data services to the clients 31. However, each processor 80 can have a unique interface, data format, query syntax, security protocol, and the like. For example, the processor 80<sub>i</sub> may be configured as a web server and so responds to hypertext transfer protocol (HTTP) request packets and generates responses formatted as hypertext markup language (HTML) web pages. In contrast other processors 80<sub>i</sub> may use file transfer protocol (FTP) or other available transfer protocols, both public and private.

In a particular example, each of the processors 80<sub>i</sub> - 80<sub>3</sub> implements a database web server that hosts a "website" having a predefined structure and access syntax and protocol. In general each website is designed to provide information about its database content response to queries (job requests) communicated through network 20. Because each processor 80<sub>i</sub> may be created and maintained by a separate organization, there is no uniform interface provided for

09872192-060501

interacting with the various processors 80<sub>i</sub>. For example, the databases may relate to customer lists in an e-commerce application, and the processor 80<sub>1</sub> may enable searches by last name, phone number, zip code, or other searchable fields whereas the processor 80<sub>2</sub> may enable searches only by last name. Accordingly, and in the absence of the service center 40, the job requests 32 must differ in content to evoke the desired response from each processor 80<sub>i</sub>. Also, the processor 80<sub>1</sub> may require that a query be processed through a series of pages before a response is generated. In contrast, the processor 80<sub>2</sub> may allow direct, one-page access to submit queries and generate responses. Hence, the processor 80<sub>1</sub> will require a series of communication packets to be received before the desired response is generated whereas the processor 80<sub>2</sub> is potentially accessed with a single communication packet that encapsulates the desired query data.

The service center 40 serves as an interface between the client 31 and the various processors 80<sub>i</sub>, and allows the client to submit data queries (job requests) without having any prior knowledge of the database being searched, and its associated search engines.

In determining which of the processors 80<sub>i</sub> to assign to complete a particular job request, the work flow controller 70 may poll the processors 80<sub>i</sub> that are coupled to the service center 40. As noted above, the processors 80<sub>i</sub> may be coupled directly to the service bus 41, or may be coupled indirectly through another communications bus, such as the Internet, for example. The polling may occur whenever a job ticket 61 is created by the job ticket service 60. Alternatively, the polling and corresponding information collection may occur on a periodic basis, and the work flow controller 70 may store information related to the processors 80<sub>i</sub>.

As an alternative to polling, processors 80<sub>i</sub> coupled to the service center 40 may monitor the job ticket service 60. The job ticket service 60 may periodically post, in a bulletin board fashion, for example, notices for job tickets that are available for processing. The processors 80<sub>i</sub> may then submit a bid for the tasks and processes defined in the job ticket notice. The work flow controller 70, or the separate, optional bidding service 90, may review the bids, and determine which single processor 80<sub>i</sub> or combination of processors 80<sub>i</sub> would be best suited to complete the tasks and processes defined in the job ticket notice.

The service center 40 may include several features to provide security and to control access to the job ticket 61. As discussed above, the job ticket service 60 may include a provision for branch locking. In addition, servers may be used to authorize and authenticate a processor 80<sub>i</sub> and

0007132 060501  
15507235

maintain the authorization and authentication during completion of a job request 32. The authentication server 92 receives authentication information from a processor 80 and the authorization server 94 uses the information to check authorization functionality. The authorization or access rights of the processor 80 may be carried as a part of the job ticket 61. The servers 92 and 94 may be hardware devices, but need not exist in the same hardware platform, and the servers 92 and 94 need not be tightly coupled. Alternatively, the functions of the servers 92 and 94 may be performed in programming stored in one of the components of the service center 40, such as the work flow controller 70, for example. Using the above-described features, the service center 40 may provide trusted authentication information about the processor 80 to the authorization server 94, and the authorization server 94 then performs its authority check functions.

The job ticket 61 may be signed with an industry standard public key encryption message digest (MD) signature, and may be protected by a public key encryption system. Hence, any user that has the public key may validate the job ticket 61 without having to communicate with the authentication server 92. These features reduce communication between distributed server applications. The features also allow the job ticket 61 to be passed from one processor 80 to another processor 80, maintaining security, without communicating with the service center 40.

In an alternative embodiment, the job ticket 61 holds authentication/access data, allowing controlled access within the service center 40 infrastructure. Resources may be protected by passwords and other mechanisms. Access to the job ticket 61 may be similarly protected. Furthermore, processors 80, with access authorization may have such access authorization invoked by listing the processors in the job ticket. The listing may be effectuated by recording a network address for the processors 80, for example. The network address may be incorporated in the bid information recorded in the job ticket 61.

Although the above description refers to development by the work flow controller 70, other components in the network 20 may be used to develop an overall work flow to complete the job request 32. For example, the job ticket service 60 may be used to develop the overall work flow.

As discussed above, the bidding service 90 may be used to receive bid information from processors 80, coupled to the service center 40. The processors 80, submit bids in response to posting of job ticket notices at the service center 40. In an embodiment, the job ticket notice is a separate object stored in the service center 40. In another embodiment, the job ticket 61 itself

serves the notice function. The work flow controller 70 may post the job ticket notices after receipt of the job request 32. Whether the bidding service 90 or the work flow controller 70 receives the bids, the bid evaluation and selection process may be the same.

The job ticket notice posted by the work flow controller 70 may include specific tasks or processes (branches) that must be completed to complete the job request 32 (see Figure 7). A simple job request 32 may have only one branch. More complex job requests 32, such as the job request illustrated in Figure 2 (i.e., print a brochure) may have many branches. Furthermore, some branches may be so interrelated that they can only be completed in a specific sequence, while other branches can be completed in a parallel or an overlapping fashion. This interrelationship may often be the result of one branch producing an output resource that is an input resource for one or more other branches. The job ticket notice may include descriptions of specific branches and their interrelationships in sufficient detail to allow the processors 80<sub>i</sub> to bid for completion of the branches. The job ticket notice may persist in the service center 40 for a specified time to allow the processors 80<sub>i</sub> to send bids. The time may be a set value (e.g., one hour) or may be based on a completion deadline specified in the job request 32.

The bidding service 90 may select bids 91 from the processors 80, based on set criteria. For example, the job request 32 may specify minimum performance requirements (e.g., a maximum cost and a completion deadline). The bidding service 90 may reject any bids that fail to satisfy the minimum performance requirements. Where the work flow controller 70 has established multiple branches, each such branch may include minimum performance requirements. The branch-specific performance requirements may be established by the work flow controller 70 based on overall performance requirements for the job ticket 61. A processor 80 that bids on a particular branch may be rejected by the bidding service 90 if the processor 80 fails to meet the minimum performance requirements.

If the client 31 does not specify any minimum performance requirements, the bidding service 90 may apply a standard set of criteria (e.g., an industry standard). In addition, the bid must satisfy any requirements for producing output resources. In this way, bids that are made in error, or that would otherwise likely be rejected, can be screened out. For example, a bid for printing inside pages of the brochure may indicate a one year completion date. Such a bid may be rejected, even in the absence of any specified performance requirements from the client 31.

1 In addition to submitting performance requirements, the client 31 may specify an evaluation  
2 algorithm for evaluating bids. For example, the client 31 may specify that cost is to be weighted  
3 twice as much as any other performance requirement.

4 In the absence of a client-specified evaluation algorithm, the bidding service 90 may apply  
5 a standard evaluation algorithm in order to rank bids for each branch in the work flow. The  
6 evaluation algorithm may apply weighting criteria, or may apply a default rule. For example, bids  
7 may be ranked based on a maximum score, where points are awarded for cost estimates below  
8 a maximum and for completion times below a maximum. Once the evaluation algorithm has been  
9 applied, the bidding service 90 ranks the bids for each branch. If only one processor 80 survives  
10 the process, that processor 80 may be automatically selected and assigned to the branch. If  
11 multiple processors 80, survive, the bidding service 90 may provide a list of such processors 80,  
12 to the work flow controller 70, which will then select the processors 80, to be assigned to the  
13 branches. Alternatively, the list may be provided to the client 31, and the client 31 may select the  
14 processor(s) 80, to complete the tasks defined in the work flow.

15 The work flow controller 70 may associate winning bids with corresponding branches, and  
16 may store the bid information with the job ticket 61. The stored bid information may include  
17 identification information that allows the authorization server 94 and the authentication server 92  
18 to permit access to job ticket branches or to the entire job ticket 61. Because the bid information  
19 is stored with the job ticket 61, a processor 80 may access those branches for which the processor  
20 80 is authorized access without having to communicate directly with the job ticket service 61. This  
21 feature allows the job ticket 60 to be passed from one processor 80 to another processor 80,  
22 which improves processing time and efficiency.

23 In an embodiment, the work flow controller 70 accesses control data of the job ticket 61  
24 to determine which processor(s) 80, should be assigned to the specific task identified in the job  
25 ticket. The work flow controller 70 may also identify which of the processors 80, would be able  
26 to meet the criteria specified in the control data, and may provide a list of such processors 80, to  
27 the client through the front end service 30. The client 31 may then select a processor(s) 80, from  
28 the list.

29 The job ticket service may, in an alternative embodiment, be implemented in whole or in  
30 part by storing program instructions on a computer-readable medium, such as a CD-ROM, for

0007392-060501  
10509267690

1 example. A computer processor may then implement method steps to provide the job ticket  
2 service by reading and executing the program instructions.

3 Figure 5A illustrates an exemplary job ticket 61. The job ticket 61 may include two parts.  
4 A first part includes a framework 62 and an optional client extension 64. The framework 62  
5 includes information, files and programming necessary to control tasks defined in the job ticket 61.  
6 The client extension 64 may include information related to a specific client (machine) and to a user  
7 of the machine. A second part includes a security module 67 that protects the job ticket 61 from  
8 unauthorized access.

9 The framework 62 may include a job identification (ID) 63, a service ID 65, a task section  
10 68, and control data 69. The job ID 63 includes a reference to a specific job, or content 51 that  
11 is stored in the job store 50. The job ID 63 also includes a reference to a particular job store 50  
12 that is used to store the content 51. An entity that acquires a reference to the job ticket 61 can use  
13 the job ID 63 to access the corresponding content 51. Thus, the network 20 shown in Figure 3  
14 may include multiple job stores 50, and the job ID 63 may be used to correlate the job ticket 61  
15 to a specific job store 50. The service ID 65 identifies a specific job ticket service 60 that stores  
16 the job ticket 61. For example, the network 20 may include multiple job ticket services 60 (not  
17 shown in Figure 3). The service ID 65 is used to correlate the job ticket 61 to the appropriate job  
18 ticket service 60.

19 The tasks section 68 (Figure 5B) may include branch definitions, and other information  
20 needed to control completion of the branches. The tasks section 68 may be structural so that each  
21 branch or node in a node tree is represented by one or more branches 66, in the tasks section 68.  
22 In this embodiment, each node in the node tree (e.g., the mode tree 10 of Figure 2) can have  
23 associated with the node, the description 95, resources 96, lock/unlock flag 97, and security  
24 features 99. In this way, the job ticket 61 reflects a hierarchical database structure. The control  
25 data 69 includes the specific instructions, parameters, and criteria for completing the task identified  
26 by the job ticket 61. The control data 69 may also include specific data required to complete tasks  
27 defined in the job ticket 61. The control data 69 may also be associated with each node in a node  
28 tree.



1 The security module 67 controls access to a specific job ticket. The security module 67  
2 may be implemented using standard encryption and access techniques, including public/private key  
3 infrastructures, for example.

4 The client extension 64 may contain "custom" information, such as user age, credit card  
5 number and zip code. Information provided in the client extension 64 may be protected by use of  
6 a public key signature, or similar feature. Hence, all client extension information will automatically  
7 be included in a Message Digest Protocol (MDP) and will affect the signature of the job ticket 61.  
8 With the above-describe job ticket architecture, many Internet-related security issues are  
9 addressed, including IP spoofing, time controlled sessions, job ticket alterations, varying  
10 authorization levels, and client-dependent persistent data storing.

11 The job ticket 61 shown in Figure 5A may be used to refer to a specific content 51 in the  
12 job store 50. Alternatively, multiple job tickets 61 may be used to refer a specific content 51, or  
13 one job ticket 61 may be used to refer to multiple contents 51. Thus, for example, one job ticket  
14 61 may specify a repetitive printing task to be completed on similar documents, each of which has  
15 a different content 51.

16 Using the network 20 shown in Figure 3, and the corresponding job ticket shown in Figure  
17 5A, a client 31 may request and have completed many different electronic services. For example,  
18 the client 31 may use the network 20 as an e-mail application.

19 Figure 5B shows the tasks section 68 in detail. The tasks section 68 may include one or  
20 more branch descriptors 66 that include information related to processing for that branch. A  
21 description segment 95 may define the tasks to be completed for each branch. Alternatively, the  
22 description segment 95 may provide a link, or handle, to a file that contains the branch description.  
23 The resources segment 96 lists input and output resources associated with the tasks defined for the  
24 branch. The lock/unlock flag segment 97 allows a flag to be set to lock and unlock a branch. A  
25 bid information segment 98 includes bid information gathered, for example, by the bidding service  
26 90. The bid information 98 may include detailed information such as the IP address of the  
27 processors authorized access to the branch, estimated performance information (e.g., estimated  
28 cost, delivery time), and other information. Alternatively, the bid information 98 may contain a link  
29 to another file containing the detailed bid information. The security segment 99 may indicate  
30 authorized security levels, and may be used as part of a public key/private key infrastructure.

Figure 5C illustrates an embodiment of the control data 69. The control data 69 includes a client address, which may be a machine address, such as an Internet protocol (IP) address. An expiration date/time segment may be used to terminate active status of the ticket 61. Once terminated, the ticket may be deleted from the job ticket service, and the corresponding content 51 may be de-referenced. That is, the contents may no longer be referenced by a specific job ticket 61. This feature may help eliminate stale data, and free up resources for other job requests 32 (see Figure 7). Finally, the control data 69 may include specific performance requirements, such as cost an delivery, warranty, required materials, price reductions based on quantity, and other requirements, for example.

The use of job tickets as XML objects allows clients to define databases, and to store data through the job ticket service 60 and the job store 50. The databases may be used to hold contact lists, addresses, and other personal data. The databases may also be used to store any other generic data. The databases could then be used in conjunction with a variety of e-services provided by the processors 80. For example, an e-mail processor 80 that provides e-mail services may be used in conjunction with a personal contact list to send e-mail messages, transfer electronic files, or to establish a chat room. The e-mail processor 80 may access the contact list at predefined intervals to send e-mail messages to a select group of e-mail addressees. Furthermore, because the service center 40 provides a single portal to processors 80 that are coupled to the communications network 35, the client 31 need not have any knowledge of the database structure, or the processing requirements of the processors 80.

In the specific application of the generic XML database to an e-mail service, the client 31 may have established, as a generic database, a list of e-mail contacts. The contacts database may then be stored in the job store 50 as a content file 51. A corresponding job ticket 61 may be stored at the job ticket service 61. The job ticket 61 includes control data needed to send and receive e-mail through the service center 40. Furthermore, the job ticket 61 serves as a pointer to data in the content file 51. In particular, the job ticket 61 may store XML data that is related to other data stored in the content file 51.

Alternatively, the job ticket 61 may store the contacts data. This alternative takes advantage of the fact that the job ticket 61 includes a vocabulary that can be extended to include the contact data, and that the vocabulary can be further extended to include properties for each



an e-mail box, or that an alert be provided whenever an e-mail message is sent. These delivery features may be stored in the contact database in the job ticket 61. Alternatively, the delivery features may be stored in a separate database (content file 51) in the job store 50, and the service center may retrieve information from this separate database when determining how to deliver the e-mail message. Specifically, the separate database may include a variety of users, along with the user's Internet address. By comparing the Internet address provided with the outgoing e-mail to the Internet addresses in the separate database, the service center 40 can determine desired delivery options of the addressee. This process for determining delivery options is transparent to the client 31 that originated the e-mail message. All that the client 31 needs to know is the contact information (e.g., the Internet address or a contact name).

The client 31 may use the job ticket service 60 to specify a number of performance features related to the e-mail service. For example, the client 31 may want the service center 40 to attempt a specified number of delivery attempts, and if delivery does not occur, to send a return message to the client 31 indicating non-delivery of the e-mail message.

In another application, the service center 40 may operate as a focal point for a client's messaging systems. The service center may receive and store messages in the job store 50, or within a job ticket 61. Preferably, a separate content file 51 or job ticket 61 is established for each client 31 having an account on the service center 40 so that all of the messages for a single client 31 will be stored in the same directory.

The network 20 shown in Figure 4 may include the necessary communications interfaces to support voice, e-mail, and facsimile transmissions. When a call arrives at the service center 40, a central processor (not shown in Figure 4) may read an incoming address signal from the originating user, and use the address to store the message in the appropriate job ticket 61.

In addition to handling incoming calls and storing the messages, the service center may coordinate a response from the client 31. For example, the client may request a specific response (e.g., and out of office notice) to an e-mail message from a user. The responses may be stored as part of the client's content file 51. The job ticket service 60 may access the content file 51, and download an appropriate response message. The job ticket service 60 may then use a processor 80 to deliver the response back to the originating user.

1 The job ticket service 60 may also include software to transform incoming messages (voice,  
2 facsimile) into XML files for storage in the content file 51.

3 The job ticket 61 or the content file 51 may contain the data indicating the preferences of  
4 the client 31. Thus, for example, when a facsimile message in the TIFF/F format is retrieved by the  
5 service center 40, the job ticket service 60 may ascertain from the data in job ticket 61 the  
6 preferred option of displaying the facsimile message and would generate the appropriate HTML  
7 files.

8 The service center 40 may be connected to a paging system processor 80. Upon the  
9 arrival of a new message, in addition to sending an e-mail message to the client's mailbox, the  
10 service center 40 may also activate the paging system processor 80 so that a pager (not shown)  
11 would be activated. In this manner, the client 31 could receive almost instantaneous notification that  
12 a message has arrived.

13 The service center 40 may use the data in the job ticket 61 to provide other information  
14 and alerts to the client. For example, the service center 40 may send a message to the client to  
15 indicate upcoming appointments, and special dates, such as birthdays. The client 31 may designate  
16 in the control data section of the job ticket 61 that birthday greeting be automatically sent to specific  
17 individuals whose information is stored in the job ticket 61.

18 The service center 40, when supporting client communications, may operate according an  
19 asynchronous manner. In an asynchronous mode of operation, information requested by the client  
20 31 may not be available and may have to be generated after the job request 32. The asynchronous  
21 mode of operation is preferred since fewer files are generated, thereby reducing the required  
22 amount of storage. Because the information requested by a client 31 may not be available, some  
23 anchors cannot specify the filename, such as "2.html," but will instead contain a command for the  
24 file. For instance, an anchor may be defined as <AHREF=/faxweb/clients/2496801/  
25 viewpage.cgi?FAX.sub.--NUM=1&PAGE=1&VIEW.sub.--MODE=FULL> for causing a CGI  
26 to run a viewpage program so that page 1 of facsimile message 1 will be displayed in a full size  
27 image. The CGI will generate the requested information when the information has not been  
28 generated, otherwise the CGI will retrieve the information and relay the information for transmission  
29 to the client 31.

30 The service center 40 can reliably receive voice, facsimile, and data messages for a plurality

1 of clients 31 and can receive more than one message for a client 31 at a single time. The messages  
2 are stored by the service center 40 and can be retrieved at the client's convenience at any time by  
3 connecting to the service center 40.

4 This use of the service center 40 provides a great number of benefits. The client 31 would  
5 not need a facsimile machine, voice mail system, or a machine dedicated for receiving data  
6 messages. The client 31 also need not worry about losing part of the message or violating the  
7 confidential nature of the messages. The client 31, of course, can still have a facsimile machine and  
8 dedicated computer for data messages. The service center 40, however, will permit the client 31  
9 to use the telephone company's call forwarding feature so that messages may be transferred to the  
10 service center 40 at the client's convenience, such as when the client 31 is away from the office.

11 Furthermore, the client 31 may be provided with a greater or fewer number of options in  
12 displaying or retrieving messages. The options may permit the client 31 to review or retrieve the  
13 messages in other formats. The options may also permit a client 31 to join two or messages into  
14 a single message, to delete portions of a message, or to otherwise alter the contents of the  
15 messages.

16 The service center 40 may restrict a client 31 to only certain types of messages. For  
17 example, a client 31 may want the service center 40 to store only facsimile messages in order to  
18 reduce costs of using the service center 40. In such a situation, the service center 40 would  
19 perform an additional step of checking that the type of message received for a client 31 is a type  
20 of message that the service center 40 is authorized to receive on the client's behalf. When the  
21 message is an unauthorized type of message, the service center 40 may ignore the message entirely  
22 or the service center 40 may inform the client 31 that a user attempted to send a message to the  
23 service center 40.

24 The service center 40 has been described as converting the messages into XML and  
25 transmitting the XML files to the client 31. The XML format, however, is only one format for  
26 exchanging information on the Internet and is actually only one type of a Standard Generalized  
27 Mark-Up Language. The service center 40 is therefore not limited to the XML format but may be  
28 practiced with any type of mixed media page layout language that can be used to exchange  
29 information on the Internet.

30 According to an embodiment, the service center 40 may be used as a file repository serving

00073192.060504

as an archive for a particular client 31 or group of clients 31. As described above, the service center 40 may maintain a list of all messages for a particular client 31, which is displayed to the client 31 when the client 31 accesses a mailbox. The service center 40 may store all messages, whether they are voice, facsimile, or data, for a client 31 in the database indefinitely. The service center 40 may therefore be relied upon by a client 31 to establish the authenticity of a message and the existence or absence of a particular message. Through the service center 40, a client 31 can therefore maintain an accurate record of all received email messages, facsimile messages, and data transfers.

In addition to serving as a file depository, the service center 40 may also function as a document management tool. When the service center 40 receives a message, the service center 40 updates a database (content file 51) with information on the message. This information includes the type of message, whether it is a facsimile message, voice message, or data message, the time and date at which the message was received, the size of the file, such as in bytes, the telephone number of the caller leaving the message, as well as other information, such as the number of pages of a facsimile message. Because the address (e.g., a telephone number or e-mail address) called is unique for each client 31, the information also includes the intended recipient of the message.

As noted above, the job ticket 61, in conjunction with other components of the service center 40, may also be used to create a persistent, generic object-based data structure, such as an XML database. An example of the use of a job ticket 61 for this purpose is illustrated in Figure 5D. The job ticket 61 includes a contacts list 84, which may be in the form of an XML database, or some other generic database. The contacts list 84 may include a structure with entries for business 85 and personal 86 use. The business 85 and personal 86 contacts structures may include entries of individuals 87, as shown. Each of the entries 87 may include specific properties, as defined above. In addition, or alternatively, each of the entries 87 may include links to other databases that provide additional information and properties about the individual.

While the use of the job ticket 61 as a XML database has generally been described with reference to an e-mail and messaging service, the job ticket 61 is not so limited. Any data that is capable of being stored in a database may be accesses and controlled using the job ticket 61.

The features described above, and shown in Figures 5A-5D may be replicated in another embodiment of a job ticket 61 in which all data related to a specific node or branch is located with

that node or branch. Using the example node-tree 10 shown in Figure 2, each node (branch) may include detailed information, and features such as resources, authorized processors 80<sub>i</sub>, lock/unlocked flag, bid information, branch description, and other information.

Figure 6 is a diagram of functions of the job ticket service 60. The primary functions of the job ticket service 60 are to store 73 the job tickets 61, and to provide access 75 to the job tickets 61, to users such as the client 31 and to the processors 80<sub>i</sub>. To accomplish these storage and access functions, the job ticket service 60 may create a job ticket reference 72 and a job resource reference 74. The job ticket service 60 also controls job content access 76, updates 77 the job tickets 61, as processes are completed and reported by the processors 80<sub>i</sub>, completes the job tickets 61, and reports 78 when all processes are completed for a specific job ticket 61, and provides an approval process 79 to allow a client 31 to approve completion of the tasks designated in the job ticket 61.

The job ticket reference 72 includes a specific reference to a corresponding job ticket 61<sub>i</sub>. The job ticket reference 72 may be used by the job ticket service 60 to allow one or processors 80<sub>i</sub> and clients 31<sub>i</sub> to access the job ticket 61. That is, instead of passing the job ticket 61 to a processor 80, the job ticket service 60 passes the job ticket reference 72. With the job ticket reference 72, the processor 80 may access all or a part of a job ticket 61 so that the processor 80 may complete one or more processes. Unlike conventional job ticket services, the job ticket service 60 retains the job ticket in storage 73, and only permits users (clients 31 and processors 80) to access the job ticket 61. This feature allows multiple processors 80 to simultaneously complete processes for the specific job request 32 related to the job ticket 61.

The job ticket service 60 may also create a resources reference 74, and may provide the resources reference 74 to the processors 80 and the clients 31 in a manner similar to that of the job ticket reference 72. As noted above with the description accompanying Figure 2, the resources may include physical devices and materials, and may include digital files. Use of the resources reference 74 may simplify data included in the job ticket 61.

Alternatively, information contained in the resources reference 74 may be included in the job ticket 61, or may be included in other files accessed by the clients 31, and the processors 80<sub>i</sub>.

Figure 7 is a diagram showing operation of selected functions of the job ticket service 60. As shown in Figure 7, the job ticket service 60 includes a job ticket 61<sub>i</sub>, which may be a



1 programming object such as that represented in Figure 2, and described above. The job ticket 61<sub>1</sub>  
2 is shown supplied to the job ticket service 60 by the client 31<sub>1</sub>. The client 31<sub>1</sub> may be a networked  
3 computer or similar device that is capable of transmitting the digital information representing the job  
4 ticket 61<sub>1</sub> to the job ticket service 60. To ensure the job ticket 61<sub>1</sub> arrives at the job ticket service  
5 60, the job ticket 61<sub>1</sub> may contain a reference to the job ticket service 60, such as the service ID  
6 65 illustrated in Figure 5A. The service ID 65 may include a network address of the job ticket  
7 service 60. For example, the service ID 65 may include a universal resource locator (URL) if the  
8 job ticket service 60 is an Internet web site.

9 Also shown in Figure 7 are client 31<sub>2</sub> and processors 80<sub>1</sub> - 80<sub>N</sub>. The processors 80<sub>1</sub> - 80<sub>N</sub>  
10 may include networked resources such as networked printers, electronic-commerce entities, such  
11 as Internet web sites, and "brick and mortar" entities, such local print shops that are coupled to the  
12 job ticket service 60 using the service bus 41.

13 The client 31 generates a job request 32 (content 51 and job ticket data). Using the front  
14 end service 30 (not shown in Figure 7) and the service bus 41, the client 31<sub>1</sub> sends the job ticket  
15 data to the job ticket service 60 and the content 51 (not shown in Figure 7) to the job store 50.  
16 The job ticket service 60 may pass the job ticket data to the work flow controller 70, which will  
17 create a job ticket 61. The content 51<sub>1</sub> and the job ticket 61<sub>1</sub> are related by the job ID 63. The  
18 job ID 63 also includes an identification of the job store 50, and a location within the job store 50  
19 in which the content 51<sub>1</sub> is stored. In an alternate embodiment, the content 51<sub>1</sub> may be stored at  
20 the client 31<sub>1</sub>, and may then be accessed by other users through the service bus 41 and the front  
21 end service 30.

22 The job ticket 61<sub>1</sub> specifies processes that must be completed to finish the job request 32.  
23 As noted above, Figure 2 illustrates processes required to print a brochure, including the inside  
24 pages and the cover. More than one processor 80<sub>1</sub> may be required to complete such a job  
25 request, or to complete the job request in the most cost-efficient and/or timely manner. The work  
26 flow controller 70 (not shown in Figure 7) can determine which of the processors 80<sub>1</sub>-80<sub>N</sub> should  
27 complete a specific process, and, if necessary, the order in which such processes should be  
28 completed. The work flow controller 70 may poll the various processors 80<sub>1</sub> to determine which  
29 may be used to complete the job request. The work flow controller 70 may then notify selected  
30 processors 80<sub>1</sub> that a job request has been registered with the job ticket service 60.

1 For each job ticket 61<sub>i</sub> received, the job ticket service 60 creates a reference 72<sub>i</sub> to the  
2 job ticket 61<sub>i</sub>. The processor 80<sub>i</sub> may request access to the job ticket 61 in order to complete one  
3 or more processes. In response, the job ticket service 60 provides the processor 80<sub>i</sub> with the job  
4 ticket reference 72<sub>i</sub>. The job ticket reference 72<sub>i</sub> is then used as an index to the job ticket 61<sub>i</sub>.  
5 The job ticket reference 72<sub>i</sub> may also be provided to other processors, such as the processor 80<sub>2</sub>,  
6 and to other clients, such as the client 31<sub>2</sub>. The processor 80<sub>2</sub> and the client 31<sub>2</sub> may then access  
7 the job ticket 61<sub>i</sub> at the same time as the processor 80<sub>i</sub> accesses the job ticket 61<sub>i</sub>. This  
8 simultaneous access allows different processes to be completed in parallel. In the example  
9 illustrated in Figure 2, the processor 80<sub>i</sub> may complete some or all the processes for the inside  
10 pages, and the processor 80<sub>2</sub> may complete the processes for the cover.

11 Figure 8 is a block diagram illustrating an example application of the control features of the  
12 job ticket service 60. The job ticket 61<sub>i</sub> is referenced to the job content 51<sub>i</sub> by the job ticket ID  
13 63, and information related to the job ticket 61<sub>i</sub> and the job content 51<sub>i</sub> is passed over the service  
14 bus 41. The processors 80<sub>i</sub> can access the job content 51<sub>i</sub> and the job ticket 61<sub>i</sub> using the service  
15 bus 41. In the illustrated example, the job ticket 61<sub>i</sub> refers to a job request 32 to print a brochure  
16 using the processes outlined in Figure 2. The processor 80<sub>i</sub> is designated by the work flow  
17 processor 70 to produce the inside pages of the brochure and the processor 80<sub>2</sub> is designated to  
18 produce the brochure cover. The processor 80<sub>i</sub> passes a job ticket access request to the job  
19 ticket service 60. The access request may include security information that allows the processor  
20 80<sub>i</sub> to access the job ticket 61<sub>i</sub> and the corresponding content 51<sub>i</sub> or job. In response, the job  
21 ticket service 60 provides a job ticket reference 62<sub>i</sub> that is used by the processor 80<sub>i</sub> to access  
22 the job ticket 61<sub>i</sub>. The processor 80<sub>i</sub> may use information in the job ticket 61<sub>i</sub> to access the  
23 content 51<sub>i</sub> stored in the job store 50. Since the processor 80<sub>i</sub> will produce only the inside pages,  
24 the processor 80<sub>i</sub> will not need access to all the information contained in the job ticket 61<sub>i</sub>.  
25 Furthermore, because the job ticket 61<sub>i</sub> remains in the job ticket service 60, other entities, such  
26 as the processor 80<sub>2</sub>, may continue to access the job ticket.

27 As the processor 80<sub>i</sub> completes various processes, the processor 80<sub>i</sub> may update the  
28 content 51<sub>i</sub> and the job ticket 61<sub>i</sub>. Thus, the job ticket 61<sub>i</sub> may reflect the latest status of the job  
29 request 32. The status reports may indicate when a node in the node tree 10 is completed, when  
30 an interim deadline is completed, when another processor may be used to complete a process, and

09873492 060501  
10005661-1

when all processing is complete. The status report may be included in a digital file that is used by the work flow controller 70, for example. The status report may also be included in a human-readable format, such as a pop-up window on a computer display screen. The processor 80<sub>1</sub> may receive the job ticket reference 72<sub>1</sub>, and may complete all scheduled processes, returning the job ticket reference 72<sub>1</sub> to the job ticket service 60. The processor 80<sub>1</sub> may also send a copy of the job ticket reference 72<sub>1</sub> to the processor 80<sub>2</sub>, so that the processor 80<sub>2</sub> may access the job ticket 61<sub>1</sub>, and the content 51<sub>1</sub> and produce the brochure cover.

Figure 9 is a flowchart illustrating an operation 100 of the job ticket service 60. The operation 100 is based on completing the inside pages nodes shown in Figure 2. The operation 100 may be at least partly under the control of the work flow controller 70, or some equivalent device. The operation 100 assumes that a job request 32 (job ticket data and content) have been passed to the service center 40, and that a job ticket service 60 has been created. The operation 100 begins at start block 101. In review and assign processors block 105, the work flow controller 70 determines which processors 80<sub>i</sub> are able and available to complete the job. The work flow controller 70, or the optional bidding service 90 may use polling or bidding features to make the determination. If more than one processor 80<sub>i</sub> is available, and can satisfy the requirements of the job ticket 61, the work flow controller 70 may assign one specific processor 80 to the job. Alternatively, the work flow controller 70 may provide a list of processors 80<sub>i</sub> to the client 31, and allow the client 31 to select one or more processors 80.

In request job ticket block 110, a processor 80, having been authorized access to a job ticket 61, sends an access request to the job ticket service 60 using the service bus 41. In block 115, the job ticket service verifies that the processor 80 may access the job ticket 61. Access may be controlled by a password, an identification, and a public key/private key security system, for example. In block 115, if the processor 80 is denied access, an error signal may be sent to the processor and/or the client 31, block 120.

In block 115, if access is authorized, the job ticket service 60 provides the processor 80 with a copy of the job ticket reference 72 corresponding to the job ticket 61, block 125. The job ticket reference 72 allows the processor 80 to access the job ticket at any time. By accessing the job ticket 61 at any time, the processor 80 is able to view an updated version of the job ticket 61 as changes are made to the job ticket 61 by other entities, including other processors 80.

1 In block 130, the job store 50 provides access to the job content 51 that is referenced by  
2 the job ticket 61. Only that part of the content 51 that may be needed by the processor 80 may  
3 be supplied by the job store 50. For example, if the processor 80 is only to generate the inside  
4 pages of the brochure, the job store 50 may not provide access to the content required to produce  
5 the brochure cover. After receiving the job ticket reference 72 and the content 51, the processor  
6 80 may perform one or more tasks using input resources to produce an interim or final output  
7 resource. With completion of each node in the node tree 10, the processor 80 may provide an  
8 input to the job ticket service 60 to allow modification of the job ticket 61, block 135. If the  
9 processor 80 completes all required processes, the processor 80 may provide a final status report  
10 to the job ticket service 60, block 140, along with any final modifications to the job ticket 61.

11 In block 145, the job ticket service 60 and the work flow controller 70 determine if any  
12 additional tasking may be required. If additional tasks are required, the work flow controller 70  
13 will ensure the appropriate processors 80 are assigned, and the operation returns to block 110.  
14 If no additional processes are required, the operation moves to block 150 and ends.

15 Figure 10 is a flowchart illustrating the routine 105 for developing a work flow and assigning  
16 processors to the work flow. The process starts in block 200. In block 205, the service center  
17 40 receives a job request 32. The job request 32 may specify performance requirements,  
18 resources, and other parameters, and may include content 51, or a link to the content 51. In block  
19 210, the work flow controller 70 defines a work flow to accomplish the tasks specified in the job  
20 request 32. The work flow may be represented by a node tree, such as the node tree 10 shown  
21 in Figure 2.

22 In block 230, the work flow controller 70 generates a job ticket 61 using the information  
23 provided by the job request 32, the work flow generated in block 210, and an appropriate job  
24 ticket template. The job ticket 61 is then stored in the job ticket service 60. Any content 51 may  
25 be stored in the job store 50.

26 The work flow controller 70 or the job ticket service 60 may create a job ticket notice, or  
27 other object, and may post the notice, block 250, at the service center 40 so that outside entities  
28 (e.g., the processors 80) may acquire sufficient information to bid on completion of the job ticket  
29 61, or a branch 66 of the job ticket 61. In an alternative embodiment, the job ticket 61 may be  
30 posted at the service center 40. If the job ticket 61 is posted, the job ticket 61 may include

mechanism to limit access to the job ticket or to limit access to certain portions of the job ticket 61. For example, the client extension 64 may not be accessible to the processors 80.

In block 270, the service center 40 receives bids from specific processors 80 and in block 290, the service center 40 evaluates the bids. In block 295, the service center 40 determines if the client 31 submitting the job request 32 intends to select the winning bid(s), or if the service center 40 makes the selection. If the client is to make the selections, in block 300, the service center 40 provides the bid information to the client 31. Then, in block 305, the service center 40 receives the selections from the client 31. If the service center 40 is to make the selections, in block 310, the service center 40 selects the winning bid(s). In block 315, the service center notifies the winning processors. The service center may also store the bid information with the corresponding job ticket 61. In block 320, the routine 105 ends.

Figure 11 is a flowchart illustrating the sub-routine 210 for defining a work flow. The sub-routine 210 starts in block 350. In block 355, the work flow controller 70 determines if the work flow will contain multiple branches. If the work flow will contain multiple branches, the work flow controller 70 defines the branches, block 360. In block 365, the work flow controller 70 selects a branch for which resources and processes are to be defined. In block 370, the work flow controller 70 defines input resources for a first process, or node. In block 375, the work flow controller 70 defines the tasks to be completed for the first process. In block 380, the work flow controller 70 determines the output resources of the first process. In block 385, the work flow controller 70 determines if another process is required for the work flow or branch. In no additional processes are required, the work flow controller 70 determines if another branch is to be defined, block 390. If another branch is to be defined, the work flow controller 70 selects another branch, block 365, and the sub-routine 210 continues. If another branch is not to be defined, the sub-routine 210 ends, block 395. The results of the work flow definition may be incorporated into the job ticket 61 (see Figure 10, block 230).

Figure 12 is a flow chart illustrating the sub-routine 250 of posting a job ticket notice or job ticket. The sub-routine 250 starts in block 400. In block 405, the work flow controller 70 determines if the work flow associated with the job ticket 61 includes multiple branches. If the work flow does not include multiple branches, the work flow controller posts the job ticket notice listing the single branch, block 410. If the work flow includes multiple branches, the work flow

1 controller 70 posts the job ticket notice with multiple branches, block 420. The sub-routine 250  
2 then ends.

3 Figure 13 is a flow chart illustrating the sub-routine 290 for evaluating bids. The sub-  
4 routine starts in block 440. In block 445, the bidding service 90 selects a first bid for analysis. In  
5 block 450, the bidding service 90 determines if the client 31 has supplied any evaluation criteria or  
6 requirements. If the client has not supplied evaluation requirements, the bidding service 90  
7 compares the selected bid to a set of standard, minimum performance requirements, which may be  
8 industry-standard requirements, block 455. In block 460, the bidding service 90 determines if the  
9 bid meets the minimum performance requirements. If the bid does not meet the minimum  
10 performance requirements, the bid is rejected, block 475. If the bid is rejected, the bidding service  
11 90 determines if additional bids were submitted, block 495. If additional bids were submitted, the  
12 bidding processor 90 returns to block 445 and selects the next bid for evaluation.

13 If in block 450, the client 31 has supplied performance requirements, the bidding service  
14 90 compares the selected bid to the client-supplied performance requirements, block 465. In  
15 block 470, the bidding service 90 determines if the selected bid meets the minimum criteria of the  
16 client-supplied performance requirements. If the minimum criteria are not met, the bidding service  
17 90 rejects the bid, block 475.

18 In blocks 470 and 460, if the minimum criteria are met, the bidding service 90 determines  
19 if the client 31 has supplied an evaluation algorithm. If the client 31 has not supplied an evaluation  
20 algorithm, the bidding service applies a standard evaluation algorithm, which may be an industry-  
21 standard algorithm, block 485. If the client has supplied an evaluation algorithm, the bidding service  
22 90 applies the client-supplied evaluation algorithm, block 490. The bidding service 90 may then  
23 store the results of the algorithm pending evaluation of all bids.

24 In block 495, the bidding service 90 determines if any bids remain to be evaluated. If  
25 additional bids remain, the sub-routine 290 returns to block 445, and the bidding service selects  
26 the next bid for evaluation. In block 495, if no additional bids remain for evaluation, the bidding  
27 service 90 ranks the bids, block 500. The sub-routine 290 then ends, block 505.

28 Figure 14 is a flowchart illustrating the routine 130 for providing access to a job ticket 61.  
29 The routine 130 begins in block 510. In block 515, the job ticket service 60 receives a job ticket  
30 reference 72 from a processor 80, and retrieves the corresponding job ticket 61, block 520.

1 In block 525, the job ticket service 60 compares the processor identification to processors  
2 listed in the job ticket 61 or branches 66 of the job ticket 61. The job ticket service 60 determines  
3 if the selected branches 66 are locked, block 530. If the selected branches 66 are not locked, the  
4 job ticket service 60 copies the selected branches 66 to the processor 80, block 535. In block  
5 550, the job ticket service 60 then determines if the selected branches 66 require locking. If the  
6 selected branches do not require locking, the routine 130 ends, block 560. If the selected branches  
7 66 require locking, the job ticket service 60 locks the selected branches 66, block 555. The  
8 routine 130 then ends, block 560.

9 In block 530, if the selected branches 66 are locked, the job ticket service 60 determines  
10 if the processor 80 intends to modify information in the selected branches 66, block 540. If the  
11 processor 80 will not modify the selected branches 66, the job ticket service 60 may provide an  
12 error message, block 545. If the selected branches 66 will be modified, the job ticket service 60  
13 may unlock the selected branches 66, block 547.

14 Figure 15 is a flow diagram of a method for allowing access to a job ticket 61. The method  
15 may execute as part of the routine 115 shown in Figure 9. The method starts with block 600. In  
16 block 605, the authentication server 94 receives authentication information from a processor 80  
17 and retrieves a job ticket 61 corresponding to a job ticket reference 72 possessed by the processor  
18 80. At this stage of the process, the job ticket 61 (excluding the public key signature field 67)  
19 contains two information fields, the framework 62 and the client extension 64. The framework 62  
20 contains information such as the service ID, client IP address, expiration date and time, and  
21 processor authorization, as previously described. The client extension 64 contains information such  
22 as credit card number and zip code, also previously described. The information in the job ticket  
23 61 (excluding the public key signature field 67) is then, for example, optionally hashed using, for  
24 example, MD5 protocol, and encrypted with a public key encryption system, block 610, generating  
25 a hash number, block 615. Other hashing or encryption techniques may also be used. The hash  
26 number is representative of the specific information contained in the job ticket 61. The hash number  
27 generated in block 615 is then encrypted using a standard public key encryption system, block 620.  
28 Encrypting the hash number with a private key prevents any user without knowledge of the public  
29 key from modifying the job information. In block 625, the job ticket 61 and the encrypted hash  
30 number are concatenated to generate the completed job ticket 61. Hence, the completed job ticket

61 information fields: 1) the framework 62, 2) the client extension 64, and 3) the public key  
signature (encrypted hash number) 67. The method then ends, block 630.

Figure 16 illustrates a process 700 for using the service center 40 for document management purposes, with reference to Figure 4. The process starts in block 705. A client 31 sends a search request to the service center 40 for a particular document or set of documents, block 710. The client 31 may issue this request with the job request 32 by clicking on a link, such as a link to "Search Documents," which may be presented to the client 31 by the service center 40 after the client 31 has been granted accesses to the client's mailbox. The service center 40 may present the client 31 with the option to search the document archives at other times, such as when the client 31 first attempts to access the mailbox, or when the URL received by the service center 40 points toward the document archives.

In response to this request, the service center 40 sends the client 31 a search query form at block 715 to allow the client 31 to define a desired search. The search query form may include an entry for each of the data fields in the content file 51. For example, the client 31 may input one or more names for a recipient and have the service center 40 search for all messages or files directed to just those recipients. The client 31 may also indicate the type of document, such as whether it is a facsimile, voice message or data file. The search query form also has entries for the date or time, which preferably accept ranges of times and dates, and an entry for the telephone number of the caller to the service center 40. The search query form may also include an entry for the size of the file or for the number of pages, which is relevant if the message is a facsimile message. The search query form may also include an entry for the document number, which may accept a range of document numbers, and also an entry for another field.

At block 720, the client 31 enters the search parameters in the search query form and returns the information to the service center 40. The client 31 may define the search about any one data field or may define the search about a combination of two or more data fields. For example the client 31 may define a search by designating the document type as a facsimile and the calling number as (XXX) 249-6801. Once the client 31 has finished defining the search, the client 31 then selects the "SEARCH" link shown at the bottom of the screen whereby the client would send the completed search query form to the service center 40.

At block 725, the service center receives the completed search query form and, through



105990-262300

a CGI, invokes one or more of the application programs 31 for performing the desired search for any files or messages falling within the parameters of the search. The results of the search are passed from the service center 40 through the CGI, at block 730, and returned to the client 31. The service center 40 may return the search results in the form of a listing of all files or messages contained within the search parameters.

At block 735, the client 31 selects the desired file or message from the listing of messages and files. For example, by clicking on the first listed document, the client 31 sends a request to the service center 40 for a viewing of that document and, in response, the service center 40 provides a viewing of the document according to the defined preferences of the client 31. The client 31 may receive a reduced size image of the first page, a full size image of the first page, reduced size images of all pages, or full size images of all pages of the facsimile message.

At block 735, the client 31 may also have the service center 40 save the search results. For example, the client 31 may input the name of JOE's FACSIMILES as the name for the search. By clicking on a "SAVE SEARCH AS" link, the name of the search is provided from the client 31 to the service center 40. At the service center 40, the CGI invokes an application program 31 to store the results of the search in the content file 51 under the designated name. The invoked application program preferably does not store the contents of all messages but rather stores a listing of the search results in the content file 51.

In the illustrated embodiments, the service center 40, and its sub-components, including the work flow controller 70 and the job ticket service 60, for example, may be implemented as a single, special purpose integrated circuit (e.g., an ASIC) having a main or central processor section for overall, system-level control, and separate circuits dedicated to performing various different computations, functions and other processes under control of the central processor section. Those skilled in the art will appreciate that the service center 40 may also be implemented using a plurality of separate, dedicated or programmable integrated or other electrical circuits or devices (e.g., hardwired electronic or logic circuits such as discrete element circuits, or programmable logic devices such as PLDs, PLAs, or PALs). The service center 40 may also be implemented using a suitably programmed general purpose computer, e.g., a microprocessor, microcontroller or other processor device (CPU or MPU), either alone or in conjunction with one or more peripheral (e.g., integrated circuit) data and signal processing devices. In general, any device or assembly of

